


Honing household helpers: Computer scientists improve robots' ability to plan, perform complex actions

May 26 2011, by Emily Finn



The researchers test their program on a real-life robot, which is able to decide that a can needs to be picked up, then plan and execute the actual motions necessary to lift it from a table. Photo: Melanie Gonick

Imagine a robot able to retrieve a pile of laundry from the back of a cluttered closet, deliver it to a washing machine, start the cycle and then zip off to the kitchen to start preparing dinner.

This may have been a domestic dream a half-century ago, when the fields of robotics and [artificial intelligence](#) first captured public

imagination. However, it quickly became clear that even “simple” human actions are extremely difficult to replicate in robots. Now, MIT computer scientists are tackling the problem with a hierarchical, progressive algorithm that has the potential to greatly reduce the computational cost associated with performing complex actions.

Leslie Kaelbling, the Panasonic Professor of Computer Science and Engineering, and Tomás Lozano-Pérez, the School of Engineering Professor of Teaching Excellence and co-director of MIT’s Center for Robotics, outline their approach in a paper titled “[Hierarchical Task and Motion Planning in the Now](#),” which they presented at the IEEE Conference on Robotics and Automation earlier this month in Shanghai.

Traditionally, programs that get robots to function autonomously have been split into two types: task planning and geometric motion planning. A task planner can decide that it needs to traverse the living room, but be unable to figure out a path around furniture and other obstacles. A geometric planner can figure out how to get to the phone, but not actually decide that a phone call needs to be made.

Of course, any robot that’s going to be useful around the house must have a way to integrate these two types of planning. Kaelbling and Lozano-Pérez believe that the key is to break the computationally burdensome larger goal into smaller steps, then make a detailed plan for only the first few, leaving the exact mechanisms of subsequent steps for later. “We’re introducing a hierarchy and being aggressive about breaking things up into manageable chunks,” Lozano-Pérez says. Though the idea of a hierarchy is not new, the researchers are applying an incremental breakdown to create a timeline for their “in the now” approach, in which robots follow the age-old wisdom of “one step at a time.”

The result is robots that are able to respond to environments that change

over time due to external factors as well as their own actions. These robots “do the execution interleaved with the planning,” Kaelbling says.

The trick is figuring out exactly which decisions need to be made in advance, and which can — and should — be put off until later.

Sometimes, procrastination is a good thing

Kaelbling compares this approach to the intuitive strategies humans use for complex activities. She cites flying from Boston to San Francisco as an example: You need an in-depth plan for arriving at Logan Airport on time, and perhaps you have some idea of how you will check in and board the plane. But you don’t bother to plan your path through the terminal once you arrive in San Francisco, because you probably don’t have advance knowledge of what the terminal looks like — and even if you did, the locations of obstacles such as people or baggage are bound to change in the meantime. Therefore, it would be better — necessary, even — to wait for more information.

Why shouldn’t robots use the same strategy? Until now, most robotics researchers have focused on constructing complete plans, with every step from start to finish detailed in advance before execution begins. This is a way to maximize optimality — accomplishing the goal in the fewest number of movements — and to ensure that a plan is actually achievable before initiating it.

But the researchers say that while this approach may work well in theory and in simulations, once it comes time to run the program in a robot, the computational burden and real-world variability make it impractical to consider the details of every step from the get-go. “You have to introduce an approximation to get some tractability. You have to say, ‘Whichever way this works out, I’m going to be able to deal with it,’” Lozano-Pérez says.

Their approach extends not just to task planning, but also to geometric planning: Think of the computational cost associated with building a precise map of every object in a cluttered kitchen. In Kaelbling and Lozano-Pérez’s “in the now” approach, the robot could construct a rough map of the area where it will start — say, the countertop as a place for assembling ingredients. Later on in the plan — if it becomes clear that the robot will need a detailed map of the fridge’s middle shelf, to be able to reach for a jar of pickles, for example — it will refine its model as necessary, using valuable computation power to model only those areas crucial to the task at hand.

Finding the ‘sweet spot’

Kaelbling and Lozano-Pérez’s method differs from the traditional start-to-finish approach in that it has the potential to introduce suboptimalities in behavior. For example, a robot may pick up object ‘A’ to move it to a location ‘L,’ only to arrive at L and realize another object, ‘B,’ is already there. The robot will then have to drop A and move B before re-grasping A and placing it in L. Perhaps, if the robot had been able to “think ahead” far enough to check L for obstacles before picking up A, a few extra movements could have been avoided.

But, ultimately, the robot still gets the job done. And the researchers believe sacrificing some degree of behavior optimality is worth it to be able to break an extremely complex problem into doable steps. “In computer science, the trade-offs are everything,” Kaelbling says. “What we try to find is some kind of ‘sweet spot’ ... where we’re trading efficiency of the actions in the world for computational efficiency.”

Citing the field’s traditional emphasis on optimal behavior, Lozano-Pérez adds, “We’re very consciously saying, ‘No, if you insist on optimality then it’s never going to be practical for real machines.’”

Stephen LaValle, a professor of computer science at the University of Illinois at Urbana-Champaign who was not affiliated with the work, says the approach is an attractive one. “Often in robotics, we have a tendency to be very analytical and engineering-oriented — to want to specify every detail in advance and make sure everything is going to work out and be accounted for,” he says. “[The researchers] take a more optimistic approach that we can figure out certain details later on in the pipeline,” and in doing so, reap a “benefit of efficiency of computational load.”

Looking to the future, the researchers plan to build in learning algorithms so robots will be better able to judge which steps are OK to put off, and which ones should be dealt with earlier in the process. To demonstrate this, Kaelbling returns to the travel example: “If you’re going to rent a car in San Francisco, maybe that’s something you do need to plan in advance,” she says, because putting it off might present a problem down the road — for instance, if you arrive to find the agencies have run out of rental cars.

Although “household helper” robots are an obvious — and useful — application for this kind of algorithm, the researchers say their approach could work in a number of situations, including supply depots, military operations and surveillance activities.

“So it’s not strictly about getting a [robot](#) to do stuff in your kitchen,” Kaelbling says. “Although that’s the example we like to think about — because everybody would be able to appreciate that.”

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Honing household helpers: Computer scientists improve robots' ability to plan, perform complex actions (2011, May 26) retrieved 26 April 2024 from <https://phys.org/news/2011-05-honing-household-helpers-scientists-robots.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.